

# Using Linguistic Graph Similarity to Search for Sentences in News Articles

Kim SCHOUTEN<sup>1</sup> and Flavius FRASINCAR

*Erasmus University Rotterdam  
PO Box 1738, NL-3000 DR  
Rotterdam, the Netherlands*

**Abstract.** With the volume of daily news growing to sizes too big to handle for any individual human, there is a clear need for effective search algorithms. Since traditional bag-of-words approaches are inherently limited since they ignore much of the information that is embedded in the structure of the text, we propose a linguistic approach to search called Destiny in this paper. With Destiny, sentences, both from news items and the user queries, are represented as graphs where the nodes represent the words in the sentence and the edges represent the grammatical relations between the words. The proposed algorithm is evaluated against a TF-IDF baseline using a custom corpus of user-rated sentences. Destiny significantly outperforms TF-IDF in terms of Mean Average Precision, normalized Discounted Cumulative Gain, and Spearman's Rho.

**Keywords.** text searching, syntax dependencies, sub-graph isomorphism

## 1. Introduction

Nowadays, a significant portion of our mental capacity is devoted to the gathering, filtering, and consumption of information. With many things that are considered to be newsworthy, like updates from friends, twitter messages from people we follow, news messages on websites, and the more classical form of news like articles and news items, the amount of textual data (not to mention multimedia content) has become too large to handle. Even when considering only news items like articles, the number is overwhelming. And while some people can safely ignore lots of the news items, others are obliged to keep up with all the relevant news, for example because of their job.

While smart heuristics like skimming and scanning texts is of great benefit, it can only go so far. People, like investment portfolio managers, who have to monitor the stock of a certain group of companies, have to keep track of all news concerning these companies, including industry-wide news, but also that of competitors, suppliers, and customers. Therefore, being able to intelligently search news on the Web, for example to rank or filter news items, is paramount. Although text searching is very old, especially in computer science terms, the advance of new paradigms like the Semantic Web, has opened the way for new ways of searching.

---

<sup>1</sup>Corresponding Author: Kim Schouten; E-mail: schouten@ese.eur.nl.

This paper addresses one of these new search techniques, namely the search for news sentences. Searching for specific sentences enables the user to both search across and within documents, with the algorithm pointing the user to exactly the sentences that matches his or her query. With a previous publication [1] outlining the general concept of such a method and [2] detailing its inner workings, this paper aims to provide a more detailed evaluation and results discussion.

## **2. Related Work**

The over two decades worth of Web research has yielded several approaches to Web news searching. The most widely used approach is based on computing similarity by means of vector distances (e.g., cosine similarity). All documents are represented as a vector of word occurrences, with the latter recording either whether that word is in the document or not, or the actual number of times the word occurs in the document. Often only the stemmed words are used in these vector representations. The main characteristic of these methods is their bag-of-words character, with words being completely stripped of their context. However, that simplicity also allows for efficient and fast implementations, a useful trait when trying to provide a good Web experience. In spite of its simplicity, it has shown to perform well in many scenarios, for example in news personalization [3], but also in news recommendation [4]. Being the de facto default in text searching, TF-IDF [5], arguably the most well-known algorithm in this category, has been chosen to serve as the baseline for the evaluation of the proposed algorithm.

With the advance of the Semantic Web, a move towards a more semantic way of searching has been made. This includes the use of natural language processing to extract more information from text and storing the results in a formally defined knowledge base like an ontology. An example of such a setup can be found in the Hermes News Portal [6, 7], where news items are annotated using an ontology that links lexical representations to concepts and instances. After processing the news items in this way, querying for news becomes a simple matter of selecting the ontology concepts of interest and all news items being annotated with these concepts are returned. Comparable to this is Aqualog [8], a question answering application which is similar in setup as Hermes, and SemNews [9], a news platform like Hermes using its own knowledge representation.

Unfortunately, because searching is performed in the ontology instead of the actual text, only concepts that are defined in the ontology and correctly found in the text can be returned to the user. A deeper problem however is caused by the fact that ontologies are formally specified, meaning that all information in the text first has to be translated to the logical language of the ontology. While translation always makes for a lossy transformation, in this case it is worse as the target language is known to be insufficient to represent certain natural language sentences. Barwise and Cooper [10] proved that first-order logic is inadequate for some types of sentences, and most ontologies are based on propositional or description logics which have even less expressive power.

## **3. Problem Definition**

Using the linguistic principles [11] of homophonic meaning specification and compositionality, a natural way of representing text is a graph of interconnected disambiguated

words, with the edges representing the grammatical relations between words. While this representation is not as rich in semantics as an ontology, it avoids the problems of ontology-based approaches while at the same time providing more semantics than traditional word-based approaches.

With both the news items and the query represented by graphs, the problem of searching for the query now becomes related to graph isomorphism: the algorithm needs to rank all sentence graphs in the news database according to similarity (i.e., the measure of isomorphism) with the graph that describes the user query. Since we need a measure of isomorphism instead of exact graph isomorphism, we cannot simply implement Ullmann's algorithm [12].

This approximate graph isomorphism has a much larger search space than regular graph isomorphism which already is an NP-complete problem [13]. There are however some constraints that make the problem more tractable. Because all graphs are both labeled and directed, they can be referred to as attributed relational graphs, which are easier to deal with in this regard than unlabeled or undirected graphs. Furthermore, missing edges in the query graphs are allowed to be present in the news item graph (i.e., this is related to induced graph isomorphism), a characteristic which also makes the problem easier to solve since now the algorithm only has to check for the query's edges in the news sentence graph and not the other way around.

We have chosen to use an augmented version of the backtracking algorithm described in [14] to compute the graph similarities. The original algorithm iterates through the graph, checking with each step whether adding that node or edge to the partial solution can still yield a valid final solution. Because of this check, partial solutions that are known to be incorrect can be pruned, thus limiting the search space. Because parse graphs are labeled graphs, nodes can only be matched to nodes when their labels are identical, again limiting the search space. However, this will not work when considering measures of similarity or approximate matches. Then, its backtracking behavior is essentially lost as adding a node never renders a solution invalid, only less relevant. Because of this we can only speak of a recursive algorithm in the case of approximate matching. Such a recursive algorithm would assign similarity scores to all nodes and edges in the solution graph, and the sum of all these similarity scores would be the final score for this solution.

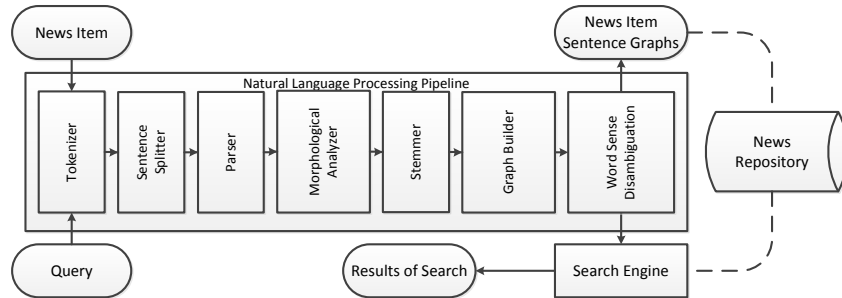
## **4. The Destiny Framework**

The Destiny framework is the implementation that follows from the above discussion. It has two main tasks: first, it transforms raw text into a graph, and second, it ranks all graphs in a database based on similarity with a given user graph. In the current use case, news items are transformed into graphs and stored in a database. The user graph represents the user query which is executed on the database.

### *4.1. News Processing*

A natural language processing pipeline has been developed that transforms raw text into a grammatical dependencies-based graph representation. The pipeline consists of a set of components with a specific natural language processing task that are consecutively

ordered, each processing the result of the previous component, sending the outcome as input to the next component in the pipeline. The same pipeline is used to process both news items and user queries. An overview of the pipeline design is given in Figure 1. The top half denotes the process of news transformation, whereas the bottom half denotes the process of searching the news.



**Figure 1.** Conceptual representation of framework [2]

The pipeline is constructed on top of the GATE framework [15]. The same framework comes packaged with an extensive set of components, hence three out of the seven components are simply standard GATE components: the tokenizer to determine word boundaries, the sentence splitter to determine sentence boundaries, and the morphological analyzer to lemmatize words. While a default GATE component exists for the Stanford Parser [16], a slightly modified version is used to take advantage of a newer version of the parser itself. Porter’s stemming algorithm [17] is used to determine the stem of each word.

The parser can be considered the main component of the pipeline, since it is responsible for finding the grammatical dependencies, thus directly influencing the graph output. Furthermore, it provides Part-of-Speech (POS) tags, essential information regarding the grammatical type of words (e.g., noun, verb, adjective, etc.). Based on the information extracted thus far, the graph builder component can construct a graph representation for all sentences. First, a node is generated for each word, encoding all known information about that word, like its POS, lemma, etc., in the node. Then, each syntactical dependency between words is used to generate an edge between the corresponding nodes, with the type of syntactical dependency encoded as an edge label. Even though a word can appear more than once in a sentence, each instantiation of that word has its own unique grammatical role in the sentence. As such it has its own dependencies, and is therefore represented as a unique node in the resulting graph as well.

An example of a graph dependencies representation of a sentence is shown in Figure 2. As can be seen, some words are integrated into an edge label, in particular prepositions and conjunctions do not receive their own node. Integrating them in an edge label gives a tighter and cleaner graph representation.

The last step of this process is the disambiguation of words, where the correct sense of a word is determined and encoded in the corresponding node. Having the word senses allows the search process to compare words, not only lexically, which would not be very accurate in a number of situations, but also semantically. Even better, the search algorithm can effectively use this information to find relations of synonymy and hypernymy

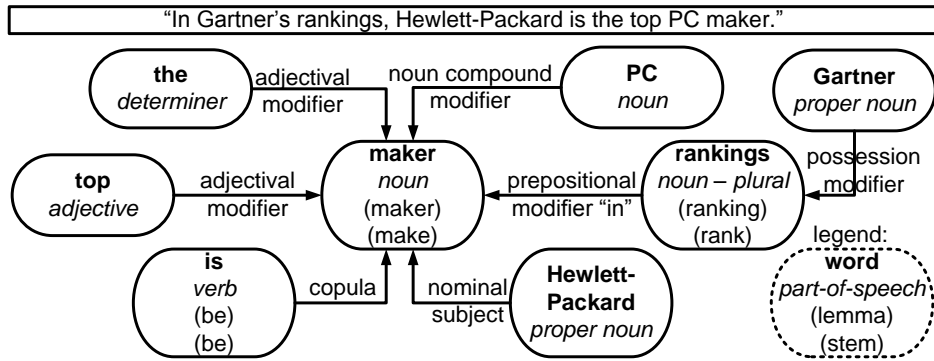


Figure 2. Graph representation of the example sentence [2]

between words, something that would not be possible otherwise. Because the development of a word sense disambiguation algorithm is outside the scope of this paper, an existing, widely used, algorithm is implemented: the simplified Lesk algorithm [18].

#### 4.2. News Searching

The news search algorithm is essentially a ranking algorithm, where all sentences in the database are ranked according to their similarity to the user query graph. As such, its core element is the part where the similarity between a sentence in the database and the user query is determined. This is the graph comparison, for which we decided to use a recursive algorithm.

However, an initial hurdle is the problem of finding a suitable starting point from where the graph comparison can commence. Since the structure of sentences can vary greatly, it would not suffice to simply start comparing at the root of both sentence graphs. On the other hand, comparing each node with every other node would be too computationally expensive. As a compromise, each noun and verb is indexed by stem and are used as starting location for the graph comparison, the intuition being that nouns and verbs are the most semantically rich words in a sentence. In practice, this means that for each noun and verb in the query sentence, an index lookup is performed, returning a list of nodes that would be suitable to start searching from for that node in the query graph. The recursive graph comparison algorithm is then executed for each of those nodes, however, each pair of (query sentence, news sentence) is associated with (and thus ranked according to) only the highest score over all runs. Suboptimal scores are discarded.

Destiny compares two graphs by first comparing the two starting nodes in the query graph and a news sentence graph. Then, using the edges of both nodes, the most suitable set of two nodes is determined to continue the graph comparison. This is done by looking one node ahead: the algorithm compares each connected node of the 'query node' with each connected node of the 'news node' to find the best possible pair. By means of a threshold, any remaining pairs with a preliminary score that is below the threshold are discarded. An additional effect of this policy is that when the preliminary score of a node is too low to be visited, its children will be discarded as well. While discarding regions of the graph that are likely to be irrelevant saves time, errors can also be introduced. As such, this is a design choice that trades off a possible increase in accuracy against a decrease in running time. The recursive process will continue until either all nodes have

been visited or no suitable matching pair is available for the remaining unvisited nodes that are connected to at least one visited node.

The similarity score of a news sentence with respect to the query sentence is essentially the sum of similarity scores of all selected pairs of nodes and pairs of edges. Thus, the actual score is determined by the similarity function of two nodes and the corresponding one for edges. While edges only have one feature, nodes have many aspects that can be compared and proper weighting of all these features can substantially improve results. As such, all feature weights are optimized using a genetic algorithm, which was described in our previous paper [1].

Nodes are compared using a stepwise comparison. First, a set of five basic features is used: stem, lemma, the full word (i.e., including affixes and suffixes), basic POS category, and detailed POS category. The basic POS category describes the grammatical word category (i.e., noun, verb, adjective, etc.), while the detailed POS category gives more information about inflections like verb tenses and nouns being singular or plural. For each feature, its weight is added to the score, if and only if the values for both nodes are identical.

If the basic POS category is the same, but the lemma's are not, there is the possibility for synonymy or hypernymy. Using the acquired word senses and WordNet [19], both nodes are first checked for synonymy and if so, the synonymy weight is added to the similarity score for this pair of nodes. If there is no synonymy, the hypernym tree of WordNet is used to find any hypernym relation between the two nodes. When such a relation is found, the weight for hypernymy, divided by the number of steps in the hypernym tree between the two nodes is added to the similarity score. In this way, very generic generalizations will not get a high score (e.g., while 'car' has a hypernym 'entity', this is so general it does not contribute much).

The last step in computing the similarity score of a node, is the adjustment with a significance factor based on the number of occurrences of the stem of that node in the full set of news items. For words which appear only once in the whole collection of news items, the significance value will be one, while the word that appears most often in the collection a significance value of zero will be assigned. Preliminary results showed that adding this significance factor, reminiscent of the inverse document frequency in TF-IDF, has a positive effect on the the obtained results. Equation 1 shows the formula used to compute the significance value for a sentence node.

$$significance_n = \frac{\log(\max\#stem) - \log(\#stem_n)}{\log(\max\#stem)} \quad (1)$$

where

- $n$  = a sentence node,
- $\#stem_n$  = how often  $stem_n$  was found in news,
- $\max\#stem$  = the highest  $\#stem$  found for any  $n$ .

#### 4.2.1. Complexity Analysis

As with any action that would require a user to wait for the results to be returned, the speed of the search algorithm is important. The query execution speed is dependent on the size of the data set, as well as on the size of the query. Furthermore, the higher the

similarity between the query and the data set, the more time it will take for the algorithm to determine how similar these two are, as the recursion will stop when encountering too much dissimilarity between the query and the current news item, as defined in the threshold parameter. To give some insight into the scalability of the algorithm with respect to the size of the data set and the size of the query, the complexity of the algorithm (in the worst case scenario) is represented in the big-O notation:

$$\mathbf{f}(n, o, p, q, r) = \mathbf{O}(no^2pqr) \quad (2)$$

where

- $n$  = the # of documents in the database,
- $o$  = the # of nodes in the query,
- $p$  = the average # of nodes in the documents in the database,
- $q$  = the # of edges in the query, and
- $r$  = the average # of edges in the documents in the database.

This can be further simplified by assuming that the number of nodes and edges in the query is equal to the average number of nodes and edges in the documents in the database, respectively. Furthermore, the average number of nodes and the average number of edges does not diverge too much, so they too can be assumed equal. In that case, the formula simplifies to two terms, as shown in Eq. 3: the number of documents in the database  $n$  and the average size of a document  $p$ .

$$\mathbf{f}(n, p) = \mathbf{O}(np^5) \quad (3)$$

Interestingly, when scaling this up, the  $p^5$  will quickly be dwarfed by  $n$ , since the size of the documents remains relatively constant, when just the number of documents increase. We can therefore conclude that the algorithm is linear in the number of documents in the database.

## 5. Evaluation

In this section, the performance of the Destiny algorithm will be measured and compared with the TF-IDF baseline. First, some insight is given into the used data set. Then the performance in terms of quality, including a discussion on the used metrics, and processing speed are given. Last, a section with advantages of using Destiny is included, as well as a failure analysis based on our experiments.

### 5.1. Setup

Since Destiny searches on a sentence level (i.e., not only among documents but also within documents), a corpus of sentences is needed where each sentence is rated against the set of query sentences. From 19 Web news items, the sentences were extracted and rated for similarity against all query sentences. The news items yielded a total of 1019 sentences that together form the data set on which Destiny will be evaluated. From this

set, ten sentences were rewritten to function as queries. The rewritten sentences still convey roughly the same meaning, but are rephrased by changing word order and using synonyms or hypernyms instead of some original words. Each sentence-query pair is rated by at least three different persons on a scale of 0 (no fit) to 3 (complete fit), resulting in a data set of over 30500 data points. For each sentence-query pair, the final user score is the average of the user ratings. From these scores, a ranking is constructed for each query of all sentences in the database. The inter-annotator agreement, computed as the standard deviation in scores is presented in Table 1. The table clearly shows that the cases where there is no fit at all are easiest to recognize. Only when a certain degree of similarity is involved does the problem become harder for humans. And even then, the more extreme cases are easier than the case of average fit where the score is between 1 and 2.

**Table 1.** Inter-annotator agreement.

| set of scores constraint                      | average standard deviation |
|---|----------------------------|
| none: all sets of scores                      | 0.1721                     |
| $0 < \text{avg}(\text{set of scores}) \leq 1$ | 0.6623                     |
| $1 < \text{avg}(\text{set of scores}) \leq 2$ | 1.0543                     |
| $2 < \text{avg}(\text{set of scores}) \leq 3$ | 0.7769                     |

As mentioned in the previous section, the weights are optimized using a genetic algorithm. The results of this optimization step are shown in Table 2, indicating that optimizing the weights is indeed a useful step in the process. Obviously, the nDCG score improves the most, since that is the optimization objective. However, the MAP also increases significantly, and while Spearman’s Rho decreases, the decrease is not statistically significant.

**Table 2.** Improvement of mean scores when optimizing.

| metric  | default | optimized | rel. improvement | t-test p-value      |
|---------|---------|-----------|------------------|---------------------|
| MAP     | 0.6806  | 0.6850    | 0.65%            | $1.768 * 10^{-8}$   |
| nDCG    | 0.1805  | 0.2112    | 16.93%           | $1.0798 * 10^{-11}$ |
| Sp. Rho | 0.2807  | 0.2766    | -1.45%           | 0.018               |

In order to have a proper evaluation, the data set is split into a training set and a test set. The split itself is made on the query level: the genetic algorithm is trained on 5 queries plus their (user-rated) results, and then tested on the remaining 5 queries. The results of the algorithm on those 5 queries are compared against the golden standard. This process is repeated 32 times, for 32 different splits of the data. All splits are balanced for the number of relevant query results, as some queries yielded a substantial amount of similar sentences, while others returned only a handful of good results.

## 5.2. Search Results Quality

The performance of Destiny is compared with a standard implementation of TF-IDF. As TF-IDF does not require training, the training set is not used and its scores are thus computed using the test set of each of the 32 splits only. The comparison is done based on

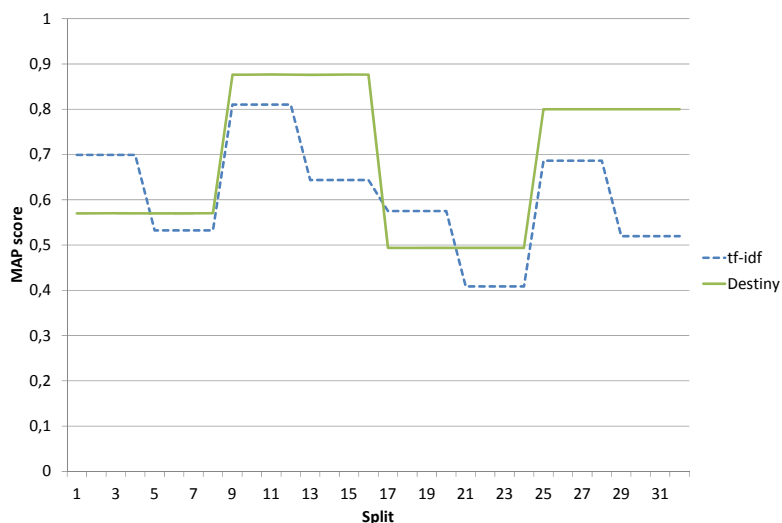


**Table 3.** Evaluation results

|         | TF-IDF mean score | Destiny mean score | rel. improvement | t-test p-value |
|---------|-------------------|--------------------|------------------|----------------|
| nDCG    | 0.238             | 0.253              | 11.2%            | < 0.001        |
| MAP     | 0.376             | 0.424              | 12.8%            | < 0.001        |
| Sp. Rho | 0.215             | 0.282              | 31.6%            | < 0.001        |

three metrics: the Mean Average Precision (MAP), Spearman’s Rho, and the normalized Discounted Cumulative Gain (nDCG) [12]. This gives a better view on the performance than when using only one metric, as each of these has its own peculiarities, as discussed in [2].

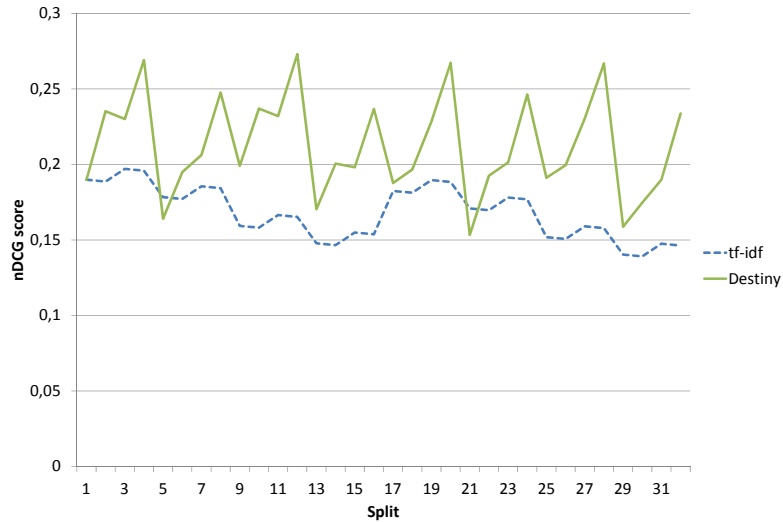
First, the results for the 32 separate training/test splits are shown for each of the three used metrics. The MAP scores are shown in Figure 3, the nDCG scores are presented in Figure 4, and the Spearman Rho numbers are given in Figure 5.

**Figure 3.** The MAP scores for Destiny and tf-idf on all splits.

To draw solid conclusions regarding the performance of the Destiny algorithm, it is compared with the TF-IDF baseline on the ranking computed from the user scores. The results, shown in Table 3, clearly show that Destiny significantly outperforms the TF-IDF baseline on the Spearman’s Rho and nDCG ranking. The p-value is computed for the paired one-sided t-test on the two sets of scores consisting of the 32 split scores for both Destiny and TF-IDF, respectively. The reported scores are the average scores over all 32 splits.

### 5.3. Processing Speed

As the system roughly consists of a news item processing step (i.e., processing raw text into a graph representation) and a graph search step, we report execution speed for these



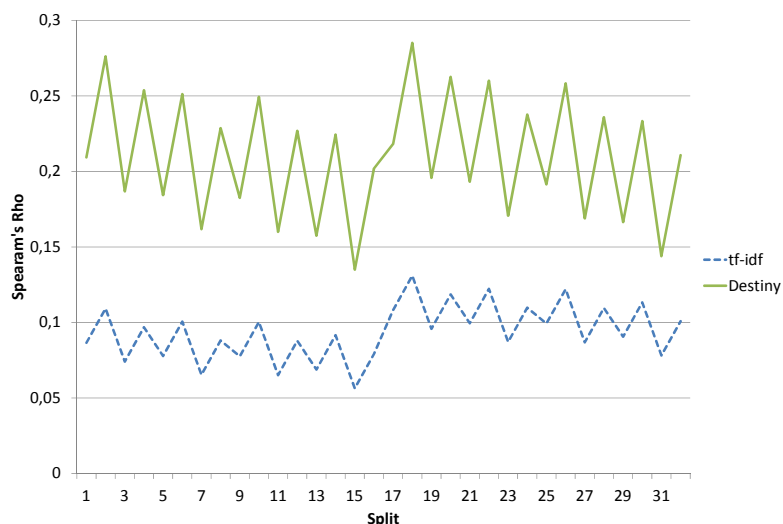
**Figure 4.** The nDCG scores for Destiny and tf-idf on all splits

two functions separately. Both were measured on an Intel Core i5 computer running at 3.8 GHz with 8 Gb RAM. In order to assess the processing speed of transforming raw text into graph representations, we let the system process a set of approximately one hundred news items and we measured the processing time per news item. The results of this experiment are reported in Figure 6. The linear pattern that can be recognized is also shown in the graph. Estimations from that linear pattern point to a processing rate of about 47 nodes per second, where the number of nodes act as a proxy for the size of the news item as the number of nodes roughly corresponds to the number of words a news item contains. Also visualized is the fact that most items are processed within 30 seconds, exceeding that amount of time only for the largest of news items (e.g., multi-page articles), of which there are only 10 in the scatter plot.

Query execution time is measured for the ten queries in our data set and compared with TF-IDF in Figure 7. The average time needed to search with Destiny is about 1570 milliseconds, while TF-IDF needs on average 800 millisecond to execute one query. As such, TF-IDF is on average approximately twice as fast as Destiny. Looking at the query execution times for Destiny, it is clearly visible how execution times vary with the size of the query and with the number of possible results. When there are clearly no similar sentences, the algorithm is fast, while if many sentence graphs need to be compared, it is slower. This influence is articulated because of the relatively small size of the set of news sentences.

#### 5.4. Advantages

Due to its focus on grammatical structure and word sense disambiguation, Destiny has some typical advantages compared to traditional search methods. The first is the focus on sentences rather than separate words. When searching is based on word occurrence



**Figure 5.** The Spearman Rho's for Destiny and tf-idf on all splits

in a document, the document can get a high score even though different words from the query are not related at all but simply occur somewhere in that document. By focusing on sentences, words from the query are at least within the same sentence, making it much more likely that they are indeed semantically related.

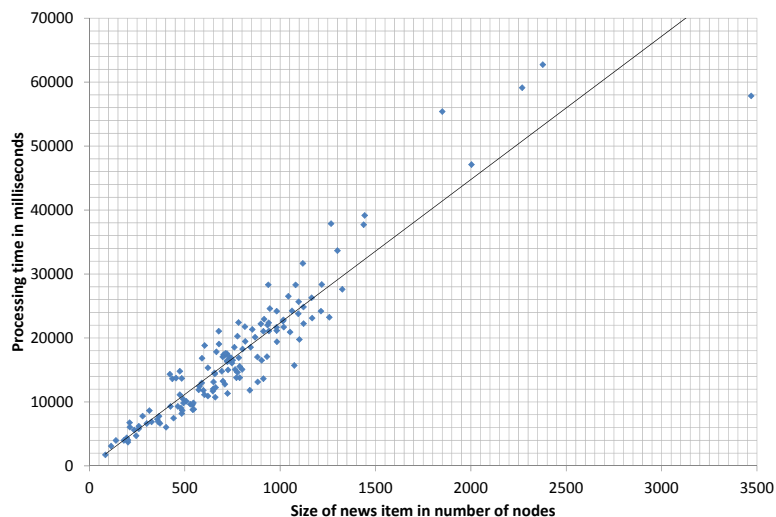
Because grammatical relations are utilized when searching, users can actively use that to search for very specific information. While many different news items can be matched to the same bag-of-words, a group of words connected by a certain grammatical structure is much more specified. As such, it is more likely that a user will find his target when he can indeed specify his search goal by means of a sentence.

While grammar can be used to specify the query, the fact that the search algorithm employs synonyms and hypernyms improves the number of hits. Using synonyms and hypernyms, sentences can be found without explicit knowledge of the words in that sentence. This is obviously a great benefit compared to traditional word-based search algorithms which only take the literal word into account.

### 5.5. Failure Analysis

In order to analyze the errors made by Destiny and assess their origin, a failure analysis has been performed. This yielded a list of situations the algorithm is not able to handle well. These situations are summarized below.

With respect to dealing with named entities, Destiny is rather limited. Various versions of a name are for example not correctly identified as being the same, neither are different names belonging to the same concept. For example, "Apple" is not recognized to be the same as "Apple Inc." or "Apple Computers Inc.", nor is it matched properly to the ticker "AAPL". Another example of the same problem would be the mismatch of the algorithm between "United States of America" and "U.S.A." or just "USA". Also,



**Figure 6.** Scatter plot of processing speed for various news item sizes

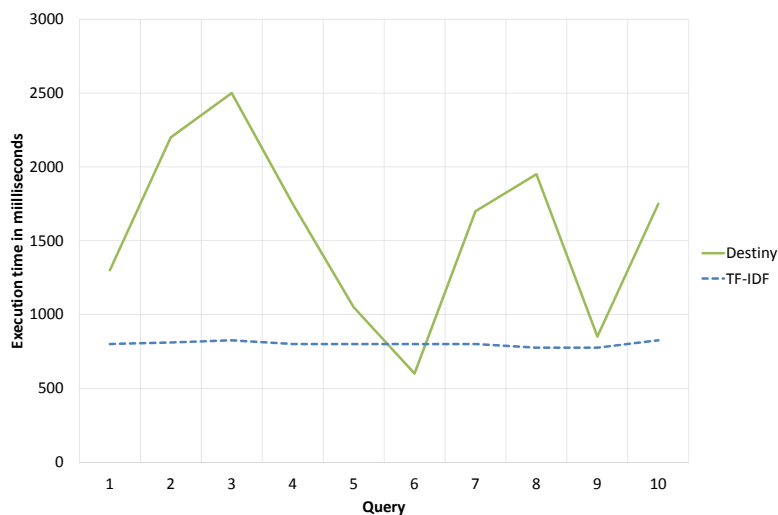
co-reference resolution is missing, so pronouns are not matched to the entity they are referring to. A graph-based approach like [11] seems particularly well suited for this work.

Also problematic in terms of semantics are proverbs, irony, and basically all types of expressions that are not to be taken literally. This caused some specific errors in the evaluation as in the data set many equivalent expressions are used for “dying”: “to pass away”, “to leave a void”, “his loss”, etc. While word synonyms can be dealt with, synonymous expressions are not considered.

Another issue is related to the fact that the search algorithm, when comparing two graphs, cannot cope well with graphs of varying size. Especially the removal or addition of a node is something the algorithm is unable to detect. When comparing Destiny with an algorithm based on graph edit distance [8], it can only detect substitution of nodes in a certain grammatical structure. Additional or missing nodes can thus break the iterative comparison, resulting in a significantly lower score than expected. For example, in the sentence “Microsoft is expanding its online corporate offerings to include a full version of Office”, it is Microsoft that is the one who will include the full version of Office, but instead of Microsoft being the grammatical subject of “include”, it is the subject of “is expanding”, which in turn is linked to “include”. When searching for “Microsoft includes Office into its online corporate offering”, a full match will therefore not be possible.

## 6. Concluding Remarks

We have shown the feasibility of searching Web news in a linguistic fashion by developing Destiny, a framework that uses natural language processing to transform both query and news items to a graph-based representation and then searches by computing the sim-



**Figure 7.** Some query execution speed measures for Destiny and TF-IDF

ilarity between the graph representing the user query and the graphs in the database. In the graph representation, much of the original semantics are preserved in the grammatical relations between the words, encoded in graph as edges. Furthermore, the search engine can also utilize semantic information with respect to words because of the word sense disambiguation component: words can be compared on a lexical level, but also on a semantic level by checking whether two words are synonyms or hypernyms.

While Destiny is slower than the TF-IDF baseline because of all the natural language processing, it is, nevertheless, better in terms of search results quality. For all three used metrics (e.g., Mean Average Precision, Spearman’s Rho, and normalized Discounted Gain), Destiny yielded a significantly higher score.

Based on the failure analysis in the previous section, it would be useful to improve the accuracy of the search results by adding a module to match named entities with different spelling or using abbreviations. Also co-reference resolution might be beneficial, as sentences later in a news item often use pronouns to refer to an entity previously introduced, while a query, being only one sentence, usually features the name of the entity. Last, as discussed in the previous section, some form of graph edit distance might be implemented to mitigate the problem of important nodes not being present in both graphs.

While not within range of real-time processing speed, the processing and query execution times of the prototype provide an acceptable basis for further development. Currently, the system is entirely single-threaded, so a multi-threaded or even distributed computing system (e.g., processing news items in parallel) is expected to improve the speed.

## Acknowledgment

The authors are partially supported by the Dutch national program COMMIT.

## References

- [1] K. Schouten and F. Frasincar. A linguistic graph-based approach for web news sentence searching. In *Proceedings of the 24th International Conference on Database and Expert Systems Applications (DEXA 2013)*, pages 57–64. Springer, 2013.
- [2] K. Schouten and F. Frasincar. Web News Sentence Searching Using Linguistic Graph Similarity. In *Proceedings of the 12th International Baltic Conference on Databases and Information Systems (DB&IS 2016)*, pages 319–333. Springer, 2016.
- [3] J. Ahn, P. Brusilovsky, J. Grady, D. He, and S. Y. Syn. Open User Profiles for Adaptive News Systems: Help or Harm? In *16th International Conference on World Wide Web (WWW 2007)*, pages 11–20. ACM, 2007.
- [4] D. Billsus and M. J. Pazzani. User Modeling for Adaptive News Access. *User Modeling and User-Adapted Interaction*, 10(2-3):147–180, 2000.
- [5] G. Salton and M.J. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, 1983.
- [6] F. Frasincar, J. Borsje, and L. Levering. A Semantic Web-Based Approach for Building Personalized News Services. *IJEBR*, 5(3):35–53, 2009.
- [7] K. Schouten, P. Ruijgrok, J. Borsje, F. Frasincar, L. Levering, and F. Hogenboom. A Semantic Web-based Approach for Personalizing News. In *ACM Symposium on Applied Computing (SAC 2010)*, pages 854–861. ACM, 2010.
- [8] V. Lopez, V. Uren, E. Motta, and M. Pasin. AquaLog: An Ontology-driven Question Answering System as an Interface to the Semantic Web. *Journal of Web Semantics*, 5(2):72–105, 2007.
- [9] A. Java, T. Finin, and S. Nirenburg. SemNews: A Semantic News Framework. In *The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference (AAAI 2006)*, pages 1939–1940. AAAI Press, 2006.
- [10] J. Barwise and R. Cooper. Generalized Quantifiers and Natural Language. *Linguistics and Philosophy*, 4:159–219, 1981.
- [11] M. Devitt and R. Hanley, editors. *The Blackwell Guide to the Philosophy of Language*. Blackwell Publishing, 2006.
- [12] J. R. Ullmann. An Algorithm for Subgraph Isomorphism. *J. ACM*, 23(1):31–42, 1976.
- [13] Stephen A. Cook. The Complexity of Theorem-proving Procedures. In *Proceedings of the third annual ACM symposium on Theory of Computing*, STOC '71, pages 151–158. ACM, 1971.
- [14] J. J. McGregor. Backtrack Search Algorithms and the Maximal Common Subgraph Problem. *Software Practice and Experience*, 12(1):23–34, 1982.
- [15] H. Cunningham, D. Maynard, K. Bontcheva, V. Tablan, N. Aswani, I Roberts, G. Gorrell, A. Funk, A. Roberts, D. Damljanovic, T. Heitz, M. A. Greenwood, H. Saggion, J. Petrak, Y Li, and W Peters. *Text Processing with GATE (Version 6)*. University of Sheffield Department of Computer Science, 2011.
- [16] D. Klein and C.D. Manning. Accurate Unlexicalized Parsing. In *41st Meeting of the Association for Computational Linguistics (ACL 2003)*, pages 423–430. ACL, 2003.
- [17] M. F. Porter. An Algorithm for Suffix Stripping. In *Readings in Information Retrieval*, pages 313–316. Morgan Kaufmann Publishers Inc., 1997.
- [18] A. Kilgariff and J. Rosenzweig. English SENSEVAL: Report and Results. In *2nd International Conference on Language Resources and Evaluation (LREC 2000)*, pages 1239–1244. ELRA, 2000.
- [19] C. Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, 1998.